IPv6 on DREN Backbone

Oct 20th, 2005

Update May 28th, 2008

Agenda

- Original implementation – Oct 20th, 2005
- Implementation requirements
- Deployment Issues
- IPv6 addressing plan summary – DREN Backbone only
- Final design
- DREN IPv6 network overview
- Status of the IPv6 implementation
- Configuration snippets for Juniper and Cisco
- Update on implementation – May 26th, 2008

IPv6 Implementation Requirements

- Requirement for IPv6 design - move away from the tunnels (ATM or GRE) and use native IPv6 in the backbone
- MPLS across the backbone
  - draft-ooms-v6ops-bgp-tunnel-00.txt
  - Juniper, “IPv6 over MPLS” – 6PE or IPv6 over Provider’s Edge
- Native IPv6 between Site Router and the SDP router
  - BGP or static routes
- Requires a new address allocation scheme using 2001:480::/32

Note – Based on 6PE, IPv6 will be over MPLS and still use IPv4 BGP sessions with the core for next-hop resolution

IPv6 Implementation - Backbone

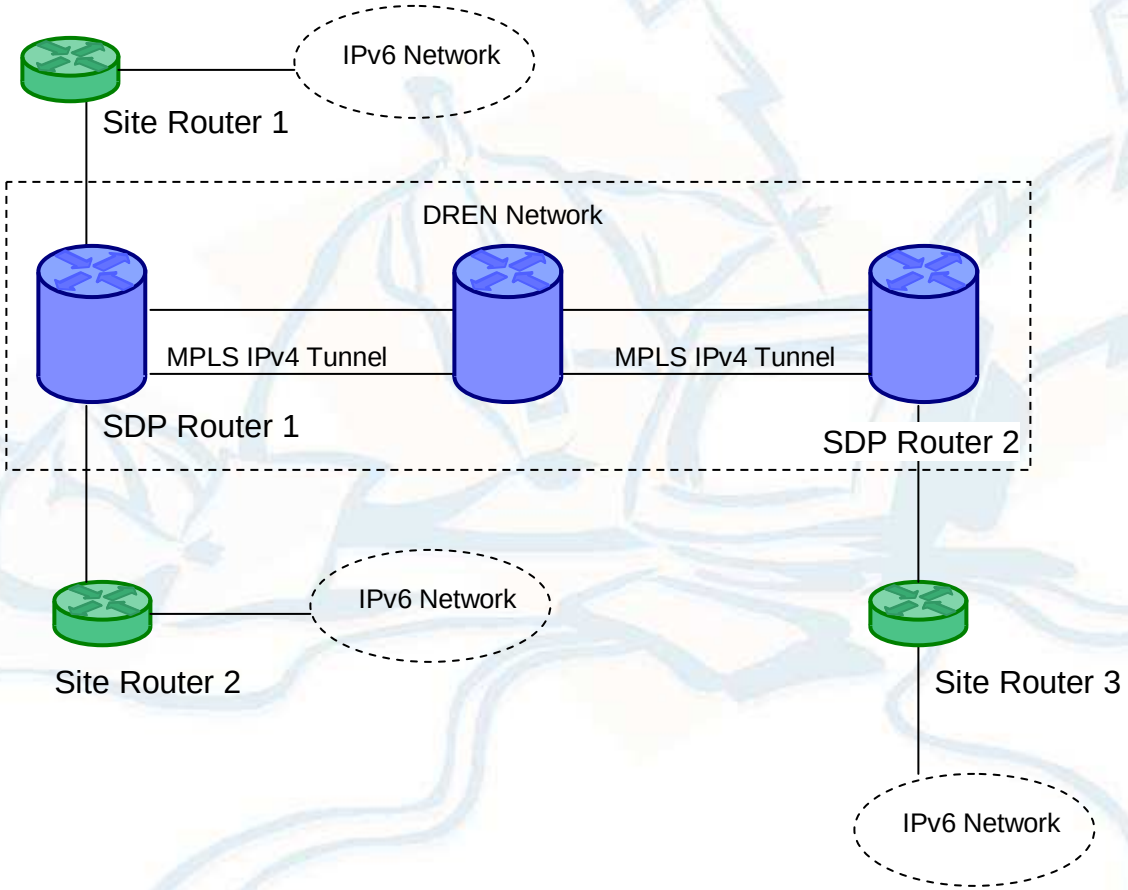

IPv6 Implementation - Edge

- Use native IPv6 between the SDP and the site router
  - If the site router will run both IPv4 and IPv6 on the same interface then only the IPv6 address needs to be added to the interfaces
  - If a separate IPv6 router will be used at the site then either use VLANs or possibly use one of the unused Fast Ethernet ports
- Configure either BGP or IPv6 static routes based on sites requirements

IPv6 Implementation - Addressing

- Original DRENv6 IPv6 allocation scheme used the old 2001:480::/35 allocation from ARIN
- New allocation from ARIN is 2001:480::/32
- Each site is given a unique /48 address space from the DREN assigned /32 address space - This is managed by DREN program office

New addressing plan discussed next

Addressing plan summary

- DREN Infrastructure - production backbone
  - 2001:0480:0000::/48
- DREN NOC – 2001:0480:0001::/48
- Reserved for future use – 2001:0480:0002::/48 – 2001:0480:000E::/48
- DRENv6 Testbed Infrastructure (going away someday)
  - 2001:0480:000F::/48
- 2001:480:0:0xxZ::/64 site-specific subnets
  - Where Z is 1 to E for site specific addresses
  - 0 is reserved for the loopback address and F is reserved for WAN interface when RFC-2547 is implemented

WAN Interface addressing

- Initially WAN interface addressing (towards vBNS+)
  - ::ffff:138.18.x.xx/128
    - where x.xx is the IPv4 address on the interface
- This mapped address is reserved for devices that do not speak IPv6 natively. According to Juniper this is an illegal implementation on the Juniper interface as the router is IPv6 capable.
- IPv6 address is removed from the WAN interfaces since IPv6 traffic is over MPLS
- WAN interface address is reserved for future RFC-2547 implementation
  - 2001:480:0:0xxF::/64 site-SDP WAN subnet

LAN interface addressing

- 2001:480:0:0xx1::/64
  - where xx is the hex version of the SDP IPv4 subnet
  - (ex. site is 138.18.88.0, so HEX(88) = 58
- SDP LAN Interface – 2001:480:0:0xx1::1/64
- Further Examples,
  - 2001:480:0:0xx0::/64 site-SDP loopback subnet
  - 2001:480:0:0xx1::/64 site-CPE1 Interface
  - 2001:480:0:0xx2::/64 site-CPE2
  - 2001:480:0:0xxF::/64 site-SDP WAN subnet
- Note: All point-to-point links use /64

IPv6 eBGP - Problem

- With the initial DREN backbone design there was an issue with the IPv6 eBGP peering – Limitation of the design
  - For IPv4 routing on the backbone, 2 loopback addresses are used on each router
    - 1st loopback is called LSP loopback and is used to create LSPs in the backbone
    - 2nd loopback is called TE loopack and is used for putting traffic over the LSPs
  - With IPv6 eBGP, the next-hop in iBGP is set as TE loopback address. This is not correct, since with 6PE, it should be LSP loopback address

IPv6 eBGP – Solution

- Option 1 - Mimic v6 design to be same as IPv4 design over LSPs.
  - Advantages – Makes design simple, scalable and consistent
  - Disadvantages – Overhaul of the backbone
    - This will require an IPv6 loopback address and a full mesh of next-hop LSP static routes on each SDP
    - This design will change with IPSec support for IPv6
- Option 2 - Leave all the IPv6 routes static and configure a default for testbed and inject in iBGP with some metrics
  - Advantages – Leaves current config as is
  - Disadvantages – Less dynamic

IPv6 eBGP – Final Design

- Mimic v6 design to be same as IPv4 design over LSPs.
  - Advantages – Makes design simple, scalable and consistent
  - Disadvantages – Overhaul of the backbone
    - This will require an IPv6 loopback address and a full mesh of next-hop LSP static routes on each SDP
    - This design will change when IPSec support for IPv6 is available
- With current architecture this design makes sense

IPv6 eBGP – Final Design – cont.

- Option 1 was decided as the final solution
- For CONUS –
  - IPv6 loopback address on each SDP and DCN
  - Full mesh of iBGP peering using v6 loopback address across all of the DREN core nodes (DCNs)
  - Each Service Delivery Point (SDP) has an IPv6 iBGP peering set up with the two closest DCNs for route-reflector configuration
  - Each SDP has a set of policies put in place to inject the customer v6 routes into iBGP
  - Each SDP has a full mesh of static routes set up to the IPV6 loopbacks of each of the other routers with a next hop of MPLS label switch path (LSP)
  - Each DCN is configured as an IPV6 Route Reflector to advertise IPV6 iBGP routes to the clients
  - Each DCN has a full-mesh IPV6 setup to exchange IPV6 routes with the other DCNs

DREN Physical Layout

- Sites physically connected to vBNS+ routers **(DS3/OCn)**
- vBNS serves as Core for carrying traffic
- DCNs Co-Located at Key vBNS+ PoPs **(Dual GigE links)**

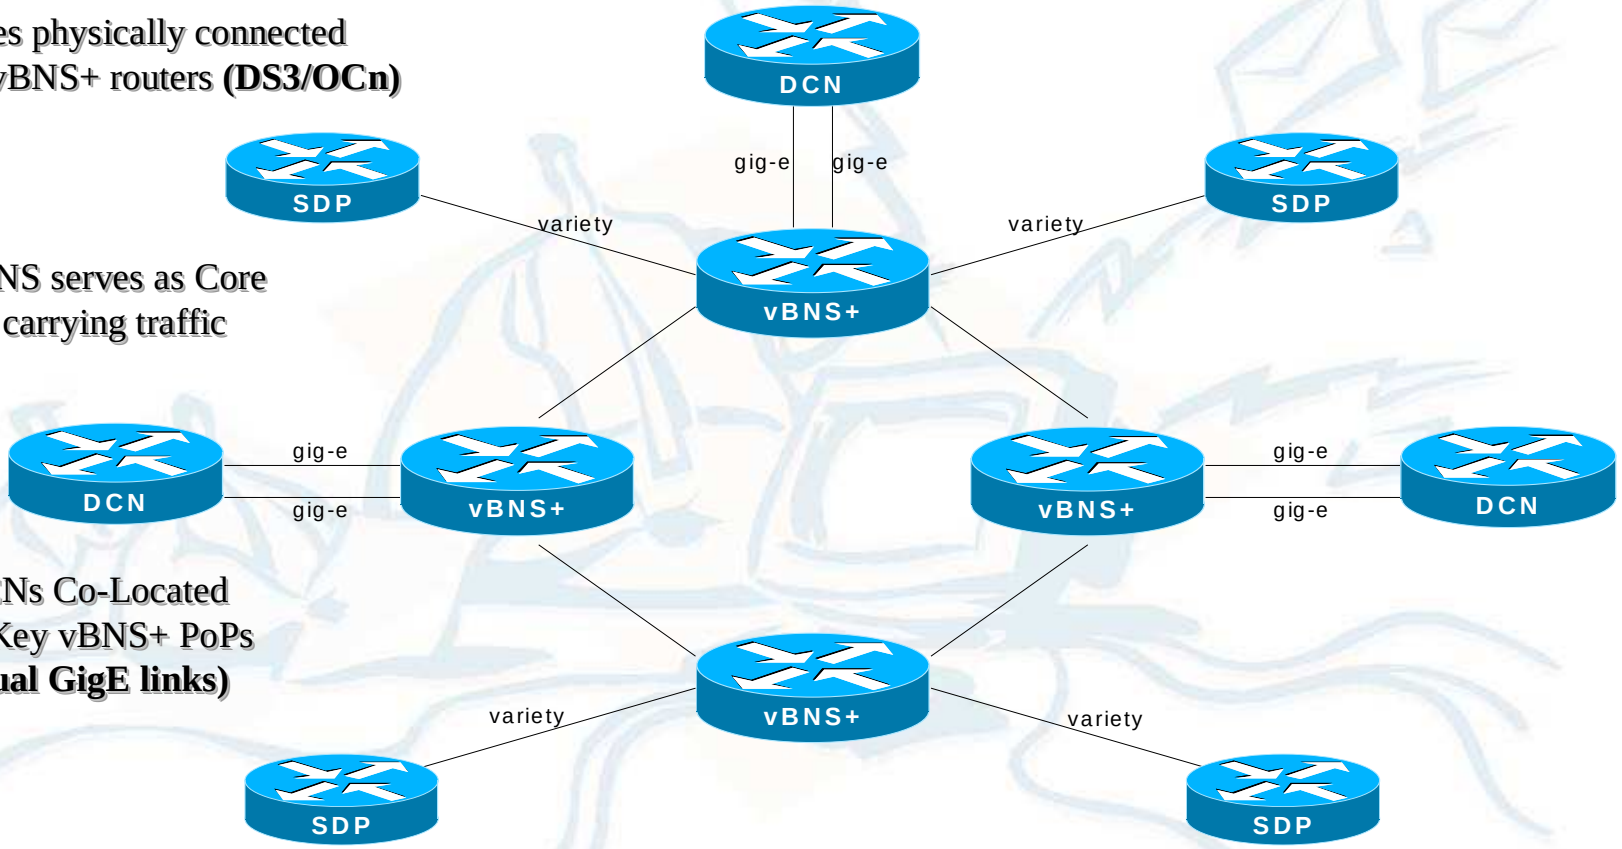

DREN Logical Layout Site Interconnectivity

- **Full Mesh of LSPs between all Sites (SDPs and DCNs)**
- **Static Routes to All SDPs over LSPs**
- **Static Routes to All DCNs over LSPs**
- All Sites are One Hop Away

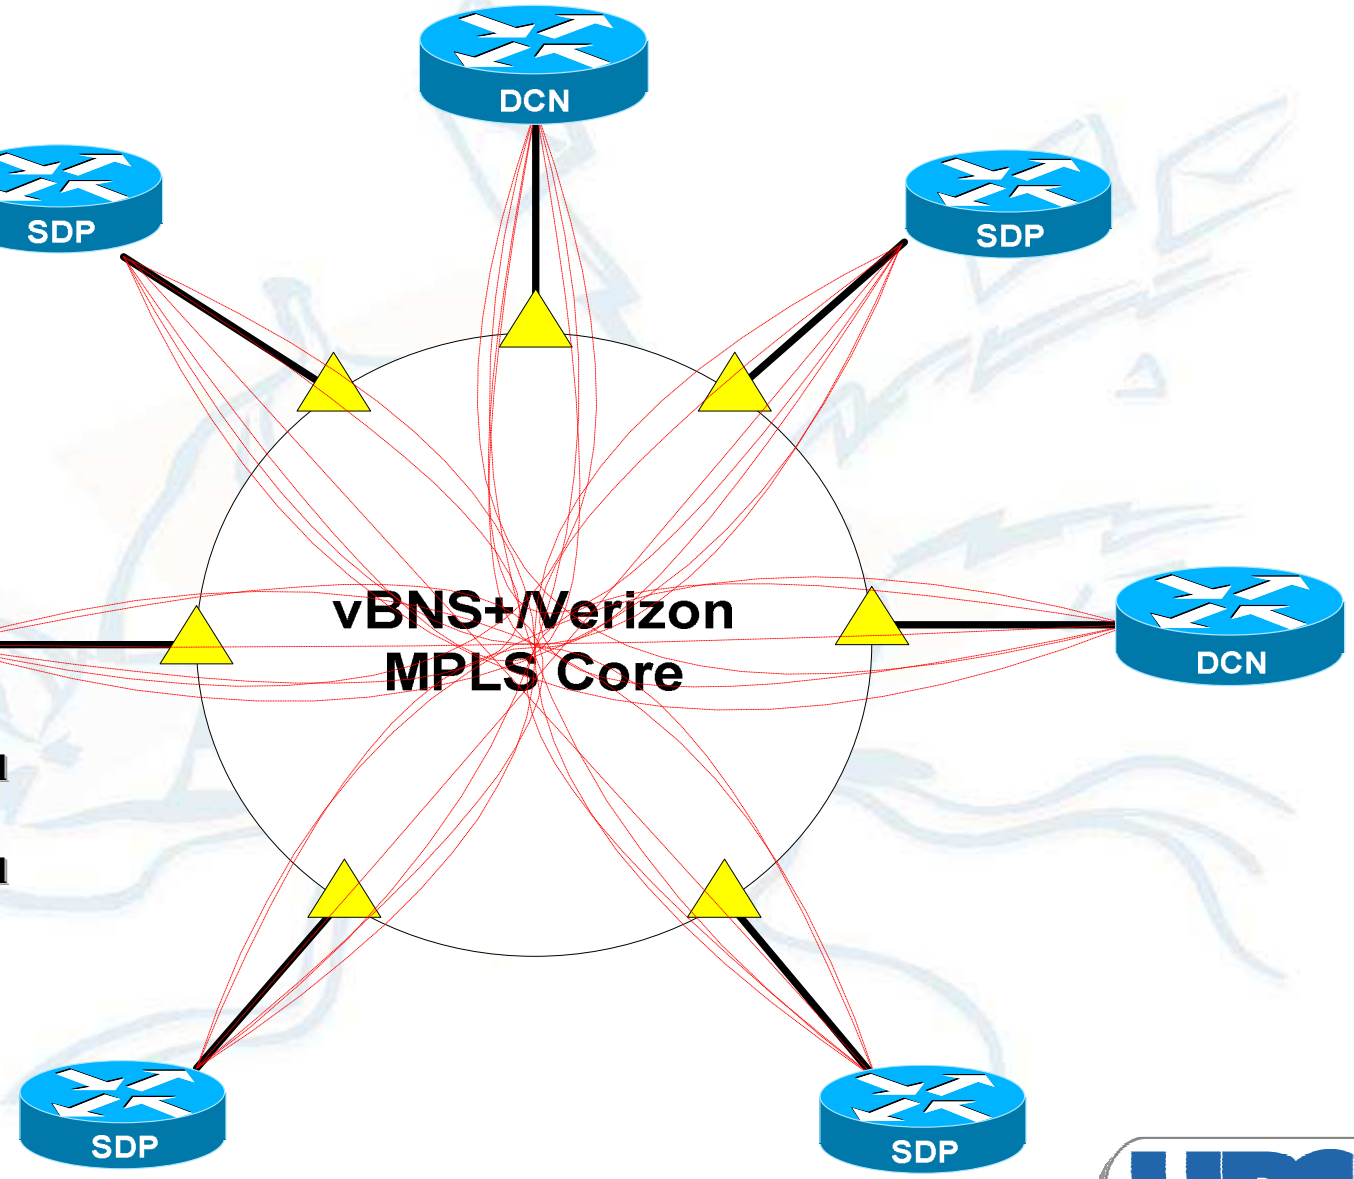

DREN Logical Layout
IBGP Route Sharing

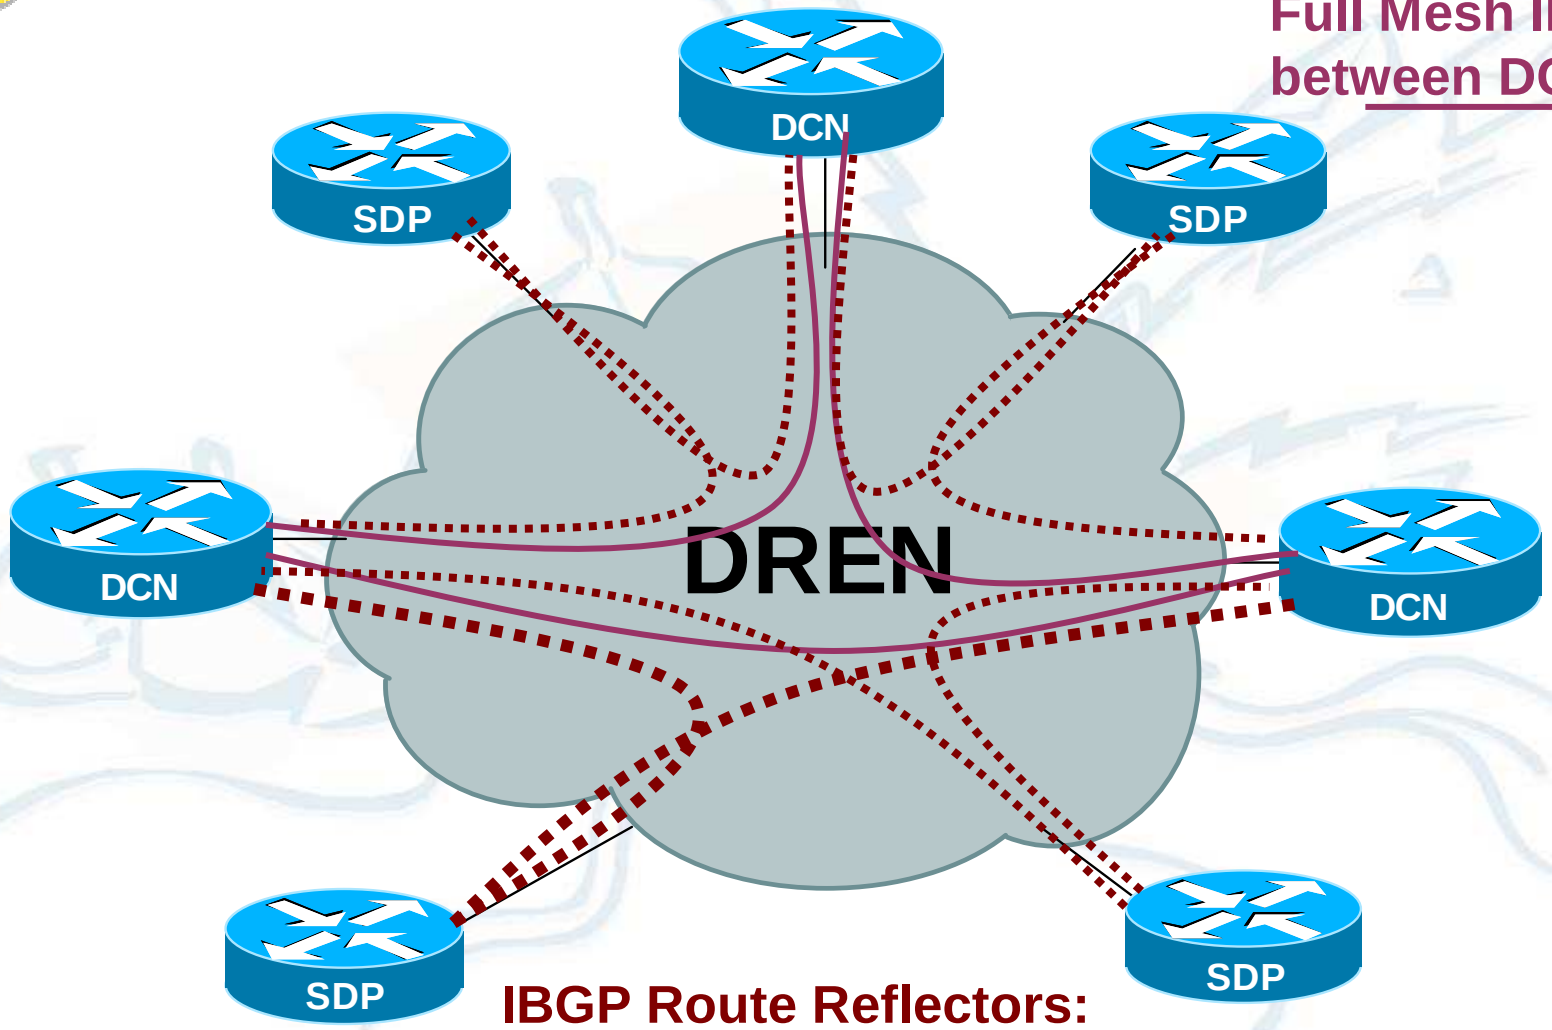

Full Mesh IBGP between DCNs

IBGP Route Reflectors: SDP to 2 Closest DCNs

DREN Logical Layout
Learning and Distributing Routes

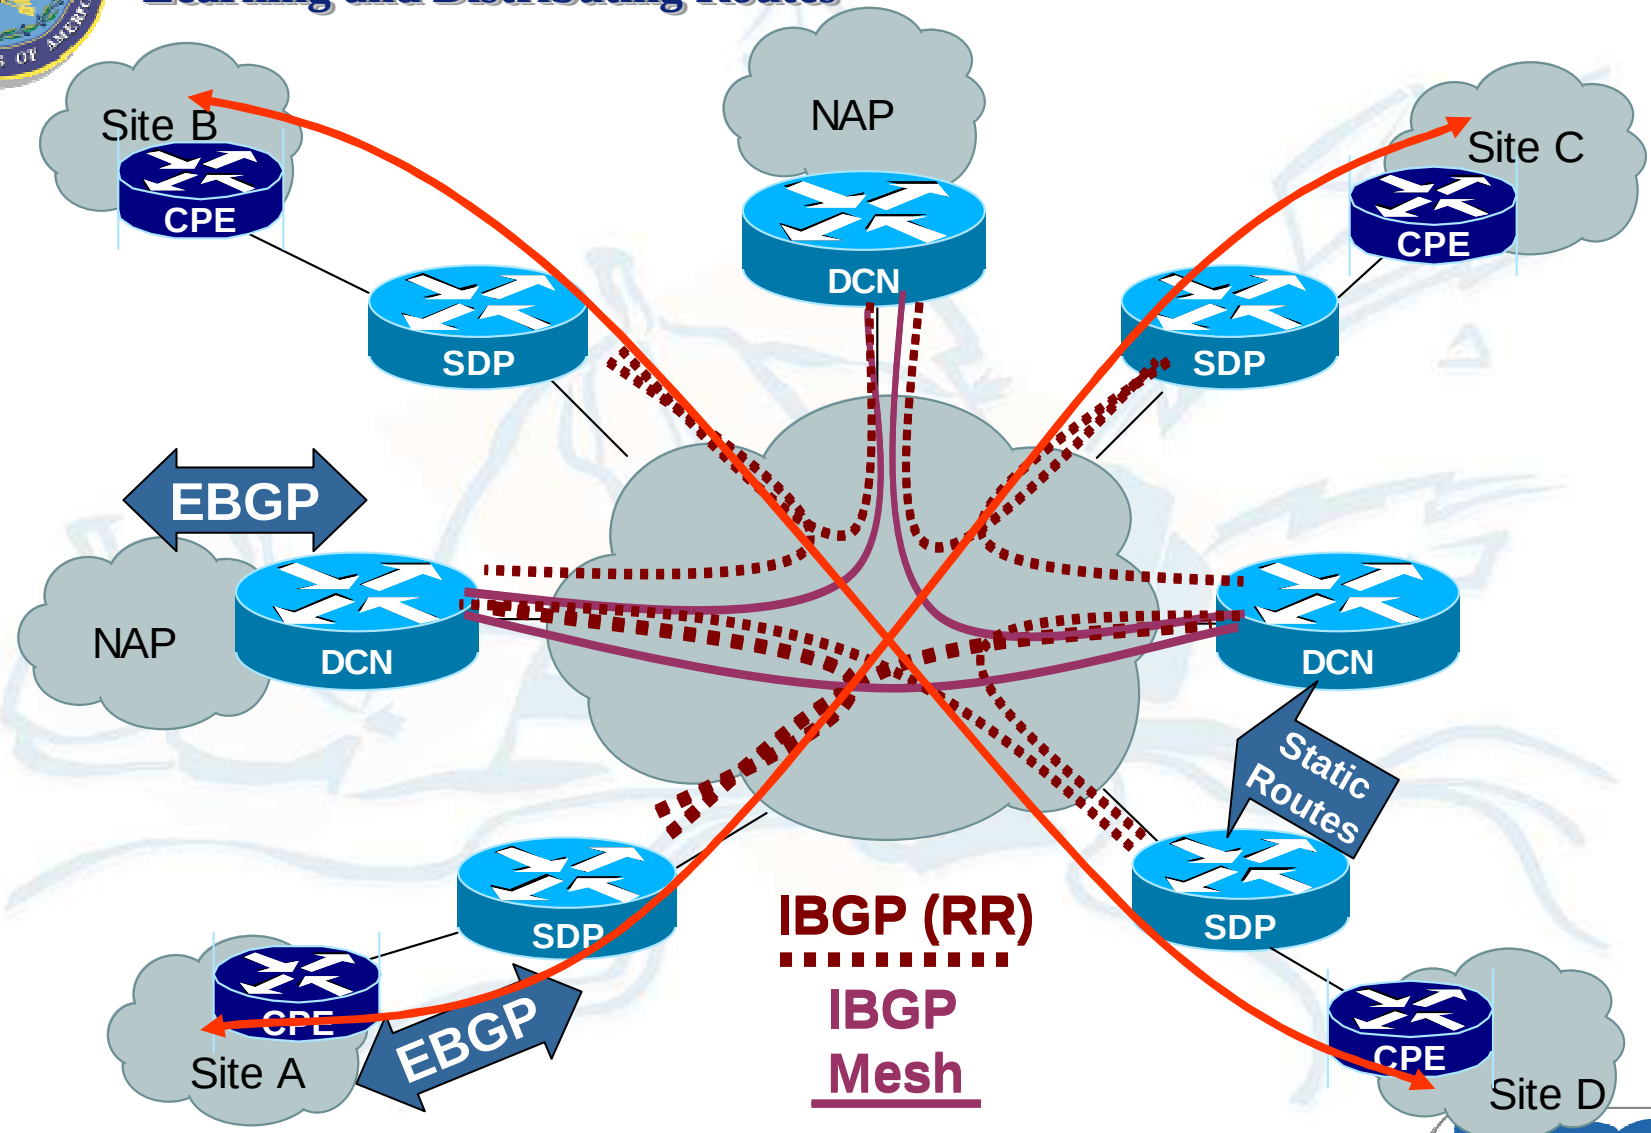

Enable v6 on Juniper Interface

- Only action required to configure v6 on an interface is to configure family inet6

```
bjones@hostname> show configuration interfaces fe-0/1/3 description "configured w/inet
    and w/inet6"; unit 0 {
   family inet {
     filter {
         input internet;
         output internet;
      }
      address 10.0.0.1/28;
      }
   family inet6 {
      address xxxx:xxx:x:xxx::x/xx;
   }
}
```

IPv6 routing options

- DREN is currently using BGP and OSPFv3
- Snippets of eBGP in the network

```
[edit protocols bgp group Ebgp-example]
 rjones@sdp1.test1# show
 type external;
 import import-from-customer;
 family inet6 {
    unicast;
 }
 export export-to-customer;
 peer-as 65043;
 neighbor XXXX:XXX:X:XXXX::X {
    local-address XXXX:XXXX:X:XXX::X;
 }
```

IPv6 routing options - Continued

- Snippets of iBGP in the network

```
[edit protocols bgp group Ibgp-example]
 rjones@sdp1.test# show
 type internal;
 family inet6 {
    unicast;
 }
 export [ send-inet6-interfaces set-customer-inet6-routes ];
 /* Route Reflector 1 */
 neighbor XXXX:XXXX:XXXX:XXXX::X;
 /* Route Reflector 2 */
 neighbor XXX:XXXX:XX:XXXX:XXX:X;
```

Enabling v6 on a Cisco router

- V6 configuration

```
Router> enable
Router# configure terminal
Router(config)# ipv6 unicast-routing
Router(config)# interface ethernet 0/0
Router(config-if)# ipv6 enable
Router(config-if)# ipv6 address 2001:0DB8:0:1::/64
Router# wr
```

Agenda

- **Update on implementation – May 26th, 2008**

IPv6 Addressing Guidance

- Sufficient space such that ease of management is usually more important than conservation of addresses
- IPv6 Working Group Conventions
  - Aggregate where possible
  - ::/48 per Enclave
  - ::/64 per interface or subnet
  - ::/128 per point to point interface
  - (::/126 for WAN or broadcast interface as required)
- Current ARIN Guidance
  - /64 when it is known that one and only one subnet is needed
  - /56 (128 subnets equivalent to an IPv4 /17) for small sites, those expected to need only a few subnets over the next 5 years.
  - /48 (256 subnets equivalent to an IPv4 /16) for larger sites

DREN IPv6 Conventions

- Follow the Working Group Guidance, except for the special use of ::/126 for interfaces or broadcast addresses are more functional than ::/64s for ease of assignment or management on the backbone ONLY.
- A ::/44 is reserved for each unclassified DREN site, allowing for 16 allocations of no less than a ::/48 for each customer (enclave) at that site.
- An additional ::/48 is reserved for all classified SDREN enclaves located at that DREN site, allowing for 16 ::/52 enclave allocations. This follows current best practice and does not create subnets with a prefix longer than ::/64.
- This plan currently does not facilitate issuing larger allocations to groups of DREN or SDREN sites that are geographically dispersed and under a single organizational structure.
- Because assignment on a nibble (4-bit) boundary is required for reverse lookup delegation, multiple delegation records may need to be configured for DREN or SDREN sites where reverse DNS will be delegated.

Customer/Enclave Allocations

- 2001:0480:0::/48; Backbone
- 2001:0480:XXX0::/44 per SDP
- 2001:0480:XXXn::/48 per Enclave (16)

  XXX = DREN SITE NUMBER 001-999
  (reserved 2001:0480:0000 – 9FFF::/48s)
  n= Enclave 0-15 (F)

Note: using only decimal XXX leaves significant unused space between allocations for growth or special cases.

IPv6 Addressing Plan (Backbone)

- 2001:0480::/32 DREN
- 2001:0480:0::/48; Backbone
- 2001:0480:0:0:0:0::/96 Loop
- 2001:0480:0:0:0:1::/96 WAN
- 2001:0480:0:rhhn::/64 LAN

- Mgmt filter trusted-sources:
  - 2001:0480::/96 (Loopbacks)
  - 2001:0480:0:0:0:1::/96 (WAN)
  - 2001:0480:0:0:0:2::/96 (IPSec)
  - 2001:0480: (NOC)

- IPSEC tunnels
  - built to loopbacks
  - 2001:0480:0:0:0:2:xxx:zzz/128 (same as 138.18.xxx.zzz of v4)

- Logical Hierarchy
- Relational Addressing (Region, Site, Interface Type, IPv4 Addres, Application, etc)
- Route Convergence more difficult on a Full Mesh
- Division of Blocks for Improved Security Posture and Granular Filter Policy